Version 1.3
November 9, 2000

US CMS COLLABORATION

US CMS Software and Computing Project Overview

1. Goals of the US CMS Software and Computing Project

The software and computing effort for CMS exceeds in scale and complexity anything that has so far been achieved in High Energy Physics. Even the new generation of experiments that will be coming on between 1999 and 2001 will not approach this scale. Because of the large number of participants and their wide geographical distribution CMS will need to employ what is essentially a new model of distributed computing and data analysis, which is without precedent in HEP. It will do so during a period of rapid change in software practices and hardware technologies. The US CMS Software and Computing Project is the response of US physicists in CMS to this challenge.

The goal of the US CMS Software and Computing Project is to provide the software and computing resources needed to enable US physicists to fully participate in the physics program of CMS. Additionally it should allow US physicists to play key roles and exert an appropriate level of leadership in all stages of the computing-related activities, from development of the reconstruction programs and software infrastructure to the extraction of physics results. This capability should extend to physicists working at their home institutions.

A key element in achieving this goal is to develop the software and to construct the facilities to provide an integrated environment for remote collaboration that would make possible central US roles in the data analysis. This includes:

1. providing the resources to support participation in the development of software associated with the design, calibration, commissioning, and analysis of the detectors in which US CMS members are involved;
2. providing the resources to support the participation in the development of reconstruction, simulation, and analysis frameworks and other physics applications infrastructure at a level appropriate to the number, capabilities, and experience of US CMS physicists; and
3. providing the resources and facilities for participation by US CMS physicists, especially those who wish to remain based in the US, in all analysis efforts and activities of interest to them.

The word ‘resources’ is meant to include personnel for development, operations, and support, as well as the hardware, commercial software purchases, and contracts for other services required to achieve these goals.

US CMS functions within the context of the full CMS experiment which in turn functions as an experiment of CERN. It is essential that this project stay well aligned with both the scientific goals of US CMS and with the policies and approaches of CMS and CERN. This project provides services and facilities that couple smoothly to CERN central computing, and to CMS facilities worldwide, so that US physicists can work productively whether at their home institutions or at CERN.

The US CMS Project is closely coordinated with the international CMS Software and Computing project that covers

(1) the computing aspects of the design, construction, evaluation and calibration of the CMS detector

(2) the storage, access and processing of event data

(3) event reconstruction and analysis, and

(4) the computing and remote collaborative infrastructure for the above.

This project supports US CMS efforts to do its appropriate share of the CMS work and also to solve special problems related to the geographical separation of US physicists from the site of the experiment.

CERN has stated clearly its policy that significant resources to support data analysis must come from sources external to CERN. This project responds to that policy by marshaling US national resources to support the analysis activities of

US physicists on CMS. The US expects to do this in a cost effective way by leveraging the knowledge, talent, and experience with HEP computing that exists within US universities and Fermilab, which is the US CMS host institution

The US also has particular issues to deal with as a result of being separated from the experimental data and the central data repositories by an ocean and by six to nine time zones. The US CMS Collaboration is itself widely spread out, as a result of the geographical expanse of the United States. US physicists are thus particularly dependent on a well-coordinated distributed data analysis system, able to deliver data and/or analysis results reliably and with acceptably short delays across transoceanic as well as national networks. These issues coupled to the unique position of the US in computing and networking technologies make it unavoidable that the US takes a lead role in these efforts, and bears the brunt of much of the R&D work and associated costs.

The long distance to the experiment means that US physicists are particularly reliant on videoconferencing for meetings and participation in the daily activities of the experiment. A high quality remote collaborative environment is required for collaborative work on the software and data analysis, and will be an important part of this project. An extension of this work will result in the ability to operate a remote “shift taking” facility in the US, including in-depth monitoring of the detector and the data acquisition system.

No remote collaborative technology can fully compensate for the out-of-phase work cycles resulting from the six to nine hour time difference between the US and Europe. The US is thus obligated to focus some of its activities nationally and sometimes regionally within the US, in order to allow the physicists of US CMS to work efficiently. This problem will be partly solved by the presence of a “Tier 1” Regional Center at Fermilab in the US Central Time zone, use of remote collaborative technologies within the US, and the use of an hierarchical “computational data grid” that places a “Tier 2” Regional Center in each of several US regions.

The remainder of this document gives an overview on the project organization and the management structure. The Work Breakdown Structure, budget, schedule and key milestones are presented.

2. Project Organization

The US CMS Software and Computing Project is involved in the CMS core framework and infrastructure software, and hardware to support the reconstruction, simulation, and physics analysis. It does not include the development of the actual reconstruction software or software for specific physics analyses, much of which will be written by physicists. A project organization is set up to carry out or oversee the project activities.

All work required for successful completion of the US CMS Software and Computing Project is organized into a work breakdown structure (WBS). The WBS contains a complete definition of the scope for the project and forms the basis for planning, execution, and control of the project. It is extended to a sufficiently low level (currently down to level 7) to make each deliverable and its provider unique and trackable. Specifically, the WBS provides the framework for cost estimating, scheduling, and budgeting.

The project has two major subprojects and the project management at level 2 of the WBS.

1. User Facilities (UF)[1]
    - 1.1. Tier 1 Regional Center at Fermilab
    - 1.2. System and User Support
    - 1.3. Operations and Infrastructure

    1.4. Tier 2 Regional Center(s)
    1.5. Networking
    1.6. Computing and Software R&D
    1.7. CMS Detector Construction Phase Computing
    1.8. Support for Fermilab based computing
2. Core Applications Software (CAS)[2]
    2.1. Software Architecture
    2.2. Interactive Graphics and User Analysis
    2.3. Distributed Data Management and Processing
    2.4. User and Developer Support and Tools
3. Project Management and Project Office (PO)
    3.1. Baselining
    3.2. Tracking
    3.3. Reporting
    3.4. PO Support

The remainder of this section provides a brief characterization of the subprojects.

2.1 The User Facilities Subproject

The mission of the User Facilities subproject is to provide the *enabling infrastructure to permit US-CMS collaborators to fully participate in the physics program of CMS from their home sites.* This enabling infrastructure will consist of the software and hardware needed to access, analyze and understand CMS data. The major cost and personnel-requirements driver for the subproject is a centralized facility in the form of a Regional Center for US CMS at Fermilab.

Regional Centers are an essential part of the network-distributed Computing Models of the LHC experiments. An effective Computing Model balances proximity to the site where the raw data is stored and first processed (CERN) against proximity to the physicists doing the analysis in each world region. For the US it is particularly important that access to the relevant portions of the data is provided within the US, mitigating the dependence on a relatively expensive transatlantic network link of limited bandwidth. A joint project of all the LHC experiments, MONARC (Models Of Network Analysis at Regional Centers)[3], is being carried out with the goal of understanding how analysis will be carried out in such an environment, and defining the necessary scope of the Regional Centers. CMS and US-CMS are participating fully and leading this joint project, and are active in shaping its work. MONARC has defined several levels ("Tiers") of Regional Center, and concluded that the first level "Tier 1" (T1) Center should have roughly 20% of the capacity of CERN ("Tier 0", T0) for a single experiment. The amount of computing required outside of the T0 center is strongly tied to the amount of computing supported by the T0 center at CERN. Therefore computing for LHC is being reviewed by an international panel of experts (Hoffman review), with a goal of defining the scope of the T0 and T1 Regional Centers based on the total amount of computing required for each of the LHC experiments. The size of the T1 Regional Center at Fermilab is being defined during this process.

The Fermilab T1 center will include substantial CPU, data storage and data access facilities. It will be able to deliver portions of the data as required to other US CMS institutions through high-speed networks and will have a high bandwidth network connection to CERN. It will also include user support personnel, personnel to manage licenses and license acquisition, and personnel to contract for needed services. It will have the responsibility and personnel to develop or acquire any software that is required to carry out its production and operation activities. It will provide a variety of training and consulting services to help university physicists carry out their computing activities at their home institutions and at Fermilab. It will also provide support for many development activities during the detector construction period before data taking begins.

In addition to raw computing hardware and software, other amenities at Fermilab are needed to enable US CMS physicists. Videoconferencing, desktop support and office space are clearly necessary; Fermilab also plans a training office and to host a remote control room for CMS. Physicists are normally most efficient when encouraged to interact freely, and we must design the user facility with the aim of encouraging physicists to hold informal discussions.

Fermilab is a major CMS collaborator in its own right. Our physicists have leading responsibilities in the hadron calorimeter, muon and trigger systems, as well as overall project management responsibility for the US CMS construction project. These subsystems require computing resources for detector design, physics simulation, analysis of test beam data, and storage and dissemination of technical information. Fermilab physicists will take leading roles in writing the reconstruction software for these subsystems, together with US CMS physicists at many universities, and this will provide a natural path toward full participation in the physics analysis of data from CMS.

The User Facility subproject also includes support for Tier 2 (T2) regional centers. A Tier 2 regional center, as defined by the MONARC project, is sized at about 2-5% of the main CMS CERN facility (and therefore 10 to 20% of the capacity of a T1 center). About five T2 regional centers will serve their local US region, connecting to a hierarchical "computational data grid". The US T2 regional center sites are not yet identified, but prototype T2 centers are being proposed now. The project scope includes hardware, software, personnel and connectivity for the centers as well as R&D and prototyping. The T1 center at Fermilab will support data import/export, documentation and software distribution to these centers.

2.2 The Core Applications Software Subproject

The Core Applications Software subproject will develop software

- to provide its share of the framework and infrastructure software required to support data analysis and simulation for CMS;
- to support the design, development, modeling, optimization, and commissioning of software related to detectors being constructed by US CMS;
- tools and services for the distributed model of computing that will enable members of US CMS to carry out data analysis while they are at home in the US or resident at CERN; and
- to satisfy any specialized needs required to carry out data analysis activities of interest to members of US CMS.

The proposed CMS software solution is based on: a hierarchy of types of software, with a professionally engineered software framework; a greater reliance on software engineering to cope with the increasing complexity; and modern programming methods, languages, and tools. This implies, at least for the foreseeable future, the use of an Object Oriented (OO) paradigm and the C++ language.

The software developed in this subproject is infrastructure software supporting the detector and physics software that is being developed primarily by physicists. The CAS subproject will coordinate software engineers to advise, train, mentor, and review the work of the physics software. Where appropriate they will help to develop specific pieces of code for which an expert engineering approach is likely to have a high payback. This will ensure that the software produced will be easy to integrate into the whole system, will be efficient in its use of hardware resources, and will be maintainable and adaptable for the full life of the CMS data analysis

In the shorter term, the core software is crucial to understand the reconstruction and triggering capabilities of CMS as long as there is still time to tune the hardware designs, while in the longer term, the functionality and quality of the production software will be a major factor in ensuring the competitiveness of the physics results. This implies that a conscious level of effort is required to support these activities.

As with any other CMS subsystem, there is a significant amount of R&D work to be done as part of the Software and Computing Project. Certain aspects of the CMS Computing Model are as yet unproven, such as the feasibility of accessing a federation of network-distributed Petabyte-scale ODBMS's with sufficient performance and reliability. The R&D work, some of which has begun, must start early enough and proceed sufficiently rapidly for there to be time to design, construct, install, and test the systems which will ultimately be used when CMS first turns on.

2.3. Project Management and Project Office

The Project Management Plan (PMP)[4] of the US CMS Software and Computing Project presents the organization, systems and plan by which the project participants will manage the project. The roles and responsibilities of the upper level project management are described in the PMP. Upper level management consists of the Level 1 Project Manager and two Level 2 Project Managers for the User Facilities and the Core Application Software subprojects each, as detailed in the PMP.

Relevant formal management systems and requirements are implemented to aid in achieving the project goals and to account properly for the use of public funds. Instruments for US CMS tracking and reporting include Statement of Work (SOW) and cost reporting templates. These documents, together with the Project Plans describing the Work Breakdown Structure of the two subprojects, constitute the appendices to the US CMS Software and Computing Project Management Plan.

The US CMS Advisory Software and Computing Board (US ASCB) provides input and feedback from the US CMS collaboration to the US CMS Software and Computing project. It advises the project managers on development of the project plan and on scientific and technical policy. The composition of the board and its relationship to US CMS and CMS are described in the PMP. It is composed of six at large members elected from the US CMS collaboration, the US CMS Physics Coordinator (also elected), and 7 ex-officio members, including the chair of the US CMS Collaboration Board, the head of the Fermilab Computing Division, the CMS Project Manager for software, the Project Manager of the US CMS Construction Project, and the L1 and L2 project managers.

The Level 1 Project Manager will control changes in requirements, costs, and schedule (in consultation and agreement, as appropriate, with the PMG and the US ASCB). Detailed change control thresholds are established in three areas: technical changes, schedule changes, and cost changes. The values of these thresholds and the authority that approves a proposed change in each area at each threshold will be set when the project baseline is established.

DOE and NSF will make funds available for support of the US CMS Software and Computing Project on an annual basis. Each year the Level 1 Project Manager will review, negotiate and approve the SOW, which will include a description of the work to be performed, the requested funds, and the manpower to be assigned to that year’s activities. Also, through reviews, the projected costs will be known at WBS level 3. Funds will then be released to the institutions that are part of the US CMS Collaboration. A management reserve will be held by the Level 1 Project Manager and will be applied during the fiscal year on the basis of performance and need.

As described in the PMP, Fermilab has management oversight responsibility for the US CMS Software and Computing project and convenes a Project Management Group, which will act as a high-level change control board for the project. The project reports costs, labor, schedule, and performance data to the Fermilab director or designee and the agency project manager.

The chair of the PMG has established a standing external review committee that periodically examines and evaluates all aspects of the US CMS Software and Computing Project. A recognized expert in HEP computing chairs the committee. Its membership includes international HEP computing experts.

3. User Facilities Subproject High Level Milestones, Work Breakdown Structure and Budget

3.1 User Facilities High Level Milestones

A timeline for the establishment of the Tier 1 regional center, together with some relevant CMS milestones[5], are listed below:

- 1999-2003: **R&D Phase**
  - 2000: CMS Milestone: select regional centers
  - 2003: CMS milestone: "Turn on" a prototype functional center
- 2004-2006 **Implementation Phase**
- 2006: CMS milestone: Fully operational centers
- 2007 on: Support of **CMS Operations Phase**

3.2 User Facilities Work Breakdown Structure

The numbers at the beginning of each of the following items are the Work Breakdown Structure levels. The User Facilities Subproject is WBS Item 1 at level 2 of the project.

WBS Item 1: User Facilities Subproject of the US CMS Software and Computing Project

WBS Item 1.1: Tier 1 Regional Center at Fermilab

This item covers the capital equipment and ongoing operational requirements of the Tier 1 Regional Center. Although hardware will be bought as late as possible in order to capitalize on falling computing prices, some hardware must be bought before 2006 in order to spread the cost over several fiscal years, to provide some reduced capacity capability before 2006 for mock data challenges and test beam analysis, and to provide greater time for full system integration

Our hardware implementation plan spreads over a three year period from 2004 to 2006. We plan to buy 10% of the total hardware needs in 2004, 30% in 2005 and the remaining 60% in 2006. There will also be ongoing expenses for the required hardware support to ensure full availability of the services and systems provided.

All hardware systems are covered here including development and test systems needed to test new hardware and software, disk, CPU and mass storage systems for the Tier 1 Regional center and both media and networked based data import/export systems.

WBS Item 1.2: System and User Support

These services and software provide the fundamental infrastructure for the support and use of the Regional Center Facilities. Because we are supporting physicists doing software development, detector design and test beam work long before the final implementation of the Tier 1 Regional Center, system and user support services are necessary during all phases of the project. In addition to supporting physicists during the building and design of the detector, they provide for the efficient and effective use of the Regional Center facilities by the US CMS physicists, Tier II regional centers, and by worldwide CMS collaborators.

This item includes documentation, support of collaborative tools, training office, software development environment, the user help desk, Tier 2 support and computer security issues.

WBS Item 1.3: Operations and Infrastructure

This item covers software license maintenance, stable system and data center operations and infrastructure support.

WBS Item 1.4: Tier 2 Regional Center(s)

The world-wide CMS computing hierarchy, which includes CERN (Tier 0) and the national regional computing centers (Tier 1) has been extended to another level, called Tier 2. An extended hierarchy consisting of 5 levels is envisioned, where group computing facilities at universities comprise Tier 3 and individual desktops and other devices make up Tier 4 (Tier 4 can be thought of as an "access layer"). This hierarchy reflects a design strategy in which the size of each Tier facility is dictated by its storage, computing and networking capabilities. Eventually, all computing facilities from Tier 0 to Tier 5 are expected to be unified in a "computational data grid", in which the facilities could be accessed from any institution in the US.

Tier 2 plays roughly the same role relative to Tier 1 as Tier 1 does to Tier 0. Each Tier 2 site has approximately 20% of the capacity (computational and disk storage) of a Tier 1 center, so five centers would have approximately the same combined capacity as the Tier 1 facility. Only five Tier 2 sites are being considered at this time, though nothing in principle restricts the number.

This item includes hardware, software and personnel for Tier 2 centers which ramp up like the Tier 1 center during the years 2004-2006. Prototype Tier 2 centers are being proposed now and resources needed for them are also in this WBS item.

WBS Item 1.5: Networking

This item covers on-site network infrastructure for the Tier 1 center as well as support for offsite network access for US CMS institutions. Network infrastructure, design, monitoring, maintenance and security are all included in this item. Transatlantic and WAN networking is assumed to be provided separately.

WBS Item 1.6: Computing and Software R&D

This item covers both hardware and software R&D activities of the User Facilities Subproject up to and including commissioning of the initial prototype regional center in 2003. These include activities carried out primarily at the site of the Tier 1 Regional Center (FNAL), however there is some overlap of these activities with the Core Applications Software subproject, especially in the software development and testing area. In general this subproject concentrates on those aspects related to deployment, operations and support, while the CAS subproject concentrates on software design and development, preparation of software releases, and end user interface aspects. Tier 2 prototype R&D activities are included in the Tier 2 WBS item.

This item also covers hardware technology tracking and testing of new technologies in a CMS context, culminating in a fully functional Regional Center in 2003, comprising roughly 5% of the final computing needs for the Tier 1 Regional Center.

WBS Item 1.7: CMS Detector Construction Phase Computing

Much of the computing power needed during the construction phase will be supplied by leveraging the use of shared FNAL computing resources. There will be a need for some dedicated CMS systems, including support for test beams, data handling and analysis, certain compute intensive simulation projects, and software development and distribution servers. These systems have a modest scope compared to the eventual production systems. The main characteristic of these systems is the substantial re-use of computing methods in use at FNAL rather than the final methods developed for production use while the experiment is under construction. Effort is reserved for coping with fluid specifications for data interchange because of similar considerations at the Tier 0 and Tier 2 sites.

The type of work necessary to provide for Construction Phase computing includes: work to configure the facilities to be useable in the CMS environment; work to modify, extend and configure CMS software to use the existing facilities.; Incremental work involved in operating and maintaining shared facilities, including lifetime upgrades; providing input to the design of the final system; Specification and construction of a modest, dedicated to CMS computing facility.

WBS Item 1.8: Support for FNAL based computing

This item covers support for CMS physicists and engineers while at Fermilab, including desktop system support with a center server and backup system. Individual desktops themselves are not included in this item. Installation and support of a remote control room for use during testbeam and CMS data taking operations is also included.

3.3 User Facilities Budget and Personnel Requirements

Costing for the User Facilities subproject is detailed in a separate document[6]. The cost of the US-CMS Regional Center is comparable to either of the present CDF or DØ Run II computing projects, and the support level per physicist is similar (it will serve the US-CMS collaboration which is similar in size to the Tevatron experiments).

The scope of the regional center is taken from the preliminary results of the Hoffman Review and is consistent with CMS estimates. The hardware costs are based on the CMS Computing Technical proposal, Run II experience, CMS estimates, and present Fermilab expenditures, and is detailed in a separate note.

We have estimated our personnel requirements for the UF subproject in two ways. Firstly using a top-down approach by scaling from our Run 2 computing experience, which is serving roughly the same number of physicists as US-CMS will be. Secondly we consulted experts from Fermilab in a bottom-up approach which led to a resource loaded WBS down to level 7 [UF WBS]. Strong agreement was found between the two approaches. l

	FY	2001	2002	2003	2004	2005	2006	2007
1	**User Facility Project**							
1.1	Tier 1 Regional Center	0	0	0	13	17	20	18
1.2	System and User Support	1	2.5	3	3	5	4.5	4.5
1.3	Operations and Infrastructure	1	1	1.5	2	2.5	5	5
1.4	Tier 2 Regional Centers	3	3	4.5	4.5	6	7.5	7.5
1.5	Networking	0.5	1	2	2.5	3	3	3
1.6	Computing and Software R&D	2.5	4	4.5	1	1	1	1
1.7	Detector Construction Phase Computing	2	2	2	0.5	0	0	0
1.8	Support for FNAL based computing	0.5	1	1	3	2	2	2
	User Facilities (total FTE)	**10.5**	**14.5**	**18.5**	**29.5**	**36.5**	**43**	**41**

Table 1: Personnel Requirements for the User Facilities subproject
The Table shows the full number of staff who will be working on User Facilities tasks, excluding physicists, but including engineers, technicians, and support staff.

In Table 2 we summarize the hardware requirements for the T1 Regional Center. For much more details see ref.[6]. The table shows the evolution of installed capacity for CPU, disks, tape, and robotic mass storage for the implementation and operation phases of the Tier 1 Regional Center.

FY	2001	2002	2003	2004	2005	2006	2007(operations)
CPU (SI95)	-	-	-	16.7k	66.8k	167k	200k
Disk (TB)	-	-	-	65	260	650	900
Tape (PB)	-	-	-	0.1	0.4	1.0	1.5
Robots	-	-	-	0.5	1	1	1.33

Table 2 Total Installed capability by year, for Implementation and CMS Operations Phases of the Tier 1 Regional Center. In order that spending is roughly equal in each year, we plan to acquire 10% of the disk and CPU in 2004, 30% in 2005, and 60% in 2006. We will buy half the tape robotics in 2004, the other half in 2005 and then budget to buy an additional one every three years during operations, hence the 0.33 robots apparently purchased in 2007. Test-bed and prototype systems for 1999-2003 are not included.

From the resource requirement the budget for the User Facility subproject is estimated, following a cost projection detailed in ref. [6]. Table 3 shows a summary by year of all hardware costs for the Tier 1 and Tier 2 centers, including networking, prototype and R&D systems, licenses, etc.

FY	2001	2002	2003	2004	2005	2006	**Total**	2007
Tier 1 hardware costs				2.64	2.84	2.60	**8.09**	2.55
Total UF at Fermilab	0.55	0.79	0.77	3.22	3.38	3.07	**11.77**	3.10
Tier 2 hardware costs	0.28	0.47	0.59	1.12	1.75	1.75	**5.96**	1.25
Total UF in Million FY01$	**0.83**	**1.26**	**1.36**	**4.34**	**5.13**	**4.82**	**17.73**	**4.35**

Table 3: User Facilities Subproject costs for the R&D, implementation, and operations Phases. The hardware costs for R&D systems cover testbed and prototype systems for the Tier 1 regional center, support of ongoing US CMS simulation, software development and test beam analysis activities at Fermilab. The cost is not yet corrected for inflation.

4. Core Applications Software Subproject High Level Milestones, Work Breakdown Structure and Budget

4.1 Core Applications Software Schedule

The software schedule is defined taking into account a number of constraints: the ongoing needs of the physicists involved in the detector and trigger design optimization and construction; the research and development required, particularly for large distributed storage systems; and the availability of resources, especially manpower. The schedule reflects an iterative development strategy, which aims to provide continuously working software while simultaneously evolving into software systems with the functionality and performance ultimately required.

The CMS Software project has defined major milestones that reflect the iterative software development process in four phases. The first phase, completed at the end of 1998, is the proof of concept for the key elements of the software architecture proposed in the CMS Computing Technical Proposal. The second phase, which was recently completed, covers the development of functional prototypes of the various software modules, which are now being exercised using large samples of simulated events and test beam data. This software has been, and is being, used for detailed studies of the detector and its sensitivity to various physics processes. The third phase involves the development of the prototypes into fully functional software while the fourth phase is the preparation of the software for production and the exercising of the complete system, in conjunction with the pre-production hardware systems. There are detailed milestones associated to the database reflecting its central importance in the CMS Computing Model.

4.2 Core Application Software Work Breakdown Structure

US CMS is playing a leading role in the design and development of CMS software in terms of the overall architecture as well as the provision of specific software deliverables. The US CMS tasks of the CAS sub-project are chosen to be of special relevance to US physicists, to maximize the US impact on the CMS project as a whole, and to match ongoing interests and efforts of the US CMS groups. They are defined and carried out in close collaboration with international CMS.

The US CMS Software and Computing CAS subproject is intimately related with the International CMS Core Computing and Software activities. In addition to appropriate levels of representation on the technical and institutional boards, US CMS has major leadership roles in both the technical development and the overall planning and management of the CMS software and computing systems. David Stickland of Princeton is Deputy Project Manager for CMS Core Software and Computing. Harvey Newman of Caltech chairs the Software and Computing Board of institutional/regional representatives. For the various CMS software projects: David Stickland coordinates the ORCA reconstruction project; Lucas Taylor of Northeastern coordinates the IGUANA interactive analysis project, and Lassi Tuura of Northeastern coordinates the CAFÉ architecture evaluation and documentation activities.

US-CMS Applications Software activities are concentrated in the following areas:

- Software Architecture
- Interactive Graphics and User Analysis (IGUANA)
- Distributed Data Management and Processing
- User Support

WBS Item 2: Core Application Software Subproject of the US CMS Software and Computing Project

WBS Item 2.1: Software Architecture

It is the intention of CMS to deploy a coherent architecture for all aspects of distributed physics data processing, including simulation, higher-level triggering, reconstruction and selection, physics analysis and visualization. The overall (or "core") architecture provides tools and a framework on which different applications will be built on. This task covers the evolution of that architecture, in particular on how the various subsystems will interface to each other through the core framework, CARF. For the subsystems to be able to talk to each other through the core framework, interfaces are required for items such as:

- persistency (including persistent object navigation and browsing, giving hints on expected behavior, setting user preferences, input sources and output destinations, dealing with meta-data etc.);
- execution models (including how to configure and use action-on-demand mechanisms);
- distributed processing (including how to instruct the system about locating replicas, controlling and configuring policies, deciding on where jobs will run and on what data sets, providing feedback to the user and querying their choices); and
- visualization (including how graphical applications can plug into the system and interact with it).

The US is taking a leading role in the overall CMS architecture and in the CAFE (CMS Architecture and Forum for Evaluation) project in particular. CAFÉ covers the evaluation and documentation side of the architecture. It provides tools for the document project, analyses the problem domain, manages the documentation, documents the different aspects of the architecture and feeds input back to architecture to evolve it. The tools must produce regular releases of entire document set such that it is internally consistent, comprehensively cross-referenced and indexed, and where each document part can be traced to its source and related documents. CAFÉ also covers research into the problem domain: developing an understanding of the problems we are to solve. It covers activities such as collecting use cases and requirements, understanding constraints, understanding the working models of the physics groups and individuals, and analyzing the impact of the possible solutions on the working models.

In addition to the overall core architectures, closely related sub-architectures will be required, such as:

- **Reconstruction sub-architecture**. As more responsibilities of reconstruction are assigned to the Physics Reconstruction and Selection, PRS, groups it will be necessary for software professionals to develop efficient software architectures for physicists to insert physics reconstruction algorithms. The on-line reconstruction for triggering and calibration as well as the off-line reconstruction architectures, which are intended to be as similar as possible, are difficult working environments requiring a high performance software architecture.
- **Detector Geometry Description Sub-Architecture** This subtask that will provide an environment for creating, manipulating, and using the parameters describing the CMS detector in a consistent manner. In particular, it covers the geometrical description of the detector elements at various levels (full engineering detail, full GEANT detail, fast simulation, trigger tower geometries, etc.), associated material properties, magnetic field map, etc. The sub-system will serve a number of clients including OSCAR, Fast Simulation, ORCA, Calibration, and User Analysis Environment.
- **Visualization Architecture** This subtask, which is closely related to the IGUANA project (WBS 2.2) covers the general architecture for graphics applications, which must cooperate with the various other (sub-) architectures and frameworks in use by CMS. It also needs to consider the issues of graphics performance in the context of complex models and distributed processing and data storage resources.

WBS Item 2.2: Interactive Graphics and User Analysis (IGUANA)

The US is solely responsible for the IGUANA software project, led by Lucas Taylor of Northeastern, that addresses interactive software needs for three domains:

- graphical user interfaces (GUI's);

- interactive detector and event visualization; and
- interactive data analysis and presentation.

IGUANA aims to support a generic and coherent set of tools which enable a wide variety of applications to be developed by both experienced and non-expert CMS software in a variety of areas including off-line simulation and reconstruction, data analysis, and test beams. Tasks include the assessment of use-cases and requirements and the evaluation, integration, adaptation, verification, deployment, and support of appropriate software in the in the CMS environment.

IGUANA recently completed successfully the "Functional Prototype Software" top-level CMS milestone. IGUANA delivered a number of well-defined deliverables to CMS including a documented set of requirements, a wide range of software prototype packages, the software infrastructure required to manage, build, test, document, release, distribute, and use the software, and a verified baseline suite of technologies with which to continue development.

The software packages delivered included a number of GUI components, a functional 3D interactive event and detector display program used in conjunction with ORCA, histogram and tag analysis components, histogram plotters, and various interactive data browsers used, for example, for browsing data in Objectivity using the CARF framework.

Wherever possible, the IGUANA developers exploit architecturally compatible software from HEP, the public domain, or commercial sector and ensure that it meets the needs and is well-integrated into a coherent overall environment. Suitable software components must generally be (‘‘small” and) modular with well-defined responsibilities and interfaces, rather than (‘‘large” and) monolithic applications. Prevalent software standards (*de-facto* or *de-jure*) are adopted wherever possible. Emphasis on modularity and standards means that if a particular package needs to be replaced or is superseded by a better package, the vast majority of the software is completely unaffected.

This approach also maximally leverages resources from outside CMS thereby allowing the relatively scarce resources available within CMS to concentrate on issues of integration and deployment, and development of the intrinsically CMS-specific components, such as the event display program. Roughly equal numbers of packages used by IGUANA come from the four sources: CMS (i.e. IGUANA), HEP, public-domain, and commercial sources.

WBS Item 2.3: Distributed Data Management and Processing

The Distributed Data Management and Processing software project develops tools to support the CMS distributed computing model. The project addresses distributed computing needs in five areas:

- Distributed Task Scheduling,
- Distributed Database Management,
- Load Balancing,
- Distributed Production Tools, and
- System Simulation.

The combination of these tools should allow CMS to exploit a global grid of computing resources, allowing physics analysis to be efficiently and transparently performed at remote sites. The various pieces of the project commence over the course of several years: the System Simulation sub-project started in November of 1998, while the Load Balancing sub-project will not begin in earnest until late 2001. Wherever possible the project seeks to develop tools that are useful immediately for CMS production at existing facilities, while developing more advanced tools for use in the future.

Each of the five sections is briefly described. The Distributed Task Scheduler sub-project is developing a complete system, which can efficiently control and schedule jobs over the computing grid. Development is proceeding well. The system currently has a scheduling mechanism that allows selection of the location to which jobs are submitted based on processor type, load, and availability of dataset.

The Distributed Database Management sub-project is developing tools external to the Object Database Management System (ODBMS) that control replication and synchronization of databases over the grid as well as performance improvement and monitoring of database access. In the short term these are fairly simple tools designed to facilitate the replication of databases produced at CERN for analysis at remote sites. CMS is currently using database replication tools developed by this sub-project in ORCA production. The complexity rapidly grows as computing resources are spread globally.

The Load Balancing Sub-Project will seek to balance the computing load over the grid of regional computing centers by combining the Distributed Task Scheduler and the Distributed Database Manager with algorithms to assess the most efficient place and method to process a request. The jobs can be sent to the data or the data replicated to the jobs, depending on the load on the computing facilities and the traffic on the network.

The Distributed Production Tools sub-project is creating tools to aid CMS production in the short term. Tools have been created to automatically transfer and archive results produced remotely to the CERN Mass Storage System (MSS). These are being used to transfer CMSIM results back to CERN from Padova, Moscow, IN2P3, Caltech, and Fermilab. Tools are being developed to submit jobs in a generic way over widely varying existing remote production centers to improve the ease of use.

The final sub-project is System Simulation, which is the CMS effort to simulate our computing facilities using the MONARC simulation toolkit. This has been used to help estimate the CMS computing needs. In turn, information collected from monitoring during CMS production on the CERN computing facility has been input to improve the quality of simulation. The simulation of the spring 2000 ORCA production was able to accurately reproduced the measured performance of the computing farm in key areas.

WBS Item 2.4: User Support

This activity supports physicists in the USA who wish to develop and use CMS software in their home institutes. This task includes the development, deployment and subsequent support of the CMS software sub-systems, associated tools, and the developers' and users' environment at US institutes, together with associated training in their use.

Whereas the Applications Software sub-project concentrates on development and initial deployment and support, the User Facilities sub-project is responsible for general support of stable production software. For example, the Applications Software covers the design, development, and test deployment of systems for distributed data analysis, while the User Facilities project includes the general deployment and support for such systems once they are reasonably stable.

The Support project attempts to provide easy access to software engineers for help with development; attempts to guarantee that US users are supported with reasonable documentation, examples, and training; and attempts to ensure that US developed software is versioned, properly configured, and easily distributed. The Support project is divided into three tasks:

- Developer Support,
- User Support, and
- Software Support Tools.

Developer support aims to provide developers of detector reconstruction, physics reconstruction, and analysis software help from professional software engineers with implementing, optimizing, and porting their code. The CAS engineers are working in a wide variety of CMS software applications and serve as a strong knowledge base for off-project developers.

The Support project offers support to users in the form of documentation, examples, and training. The Support project ensures that documentation for software developed by US-CMS is current and available to users. This includes the preparation, in electronic form, of documentation for CMS specific software: ORCA, IGUANA, OSCAR, Distributed Data Management and Processing tools, etc. Support seeks to provide a set of reliable and current examples for the use of the packages necessary to perform physics analysis in CMS to enable users to more quickly become proficient in the CMS working environment. In order for US physicists working in CMS to be successful they must be trained to use and develop software in the CMS environment. This involves the participation of all CAS engineers in yearly CMS Software Training for new users on ORCA, IGUANA, CARF, the user environment, and other software packages necessary for physics success.

For software to be useful it must be available for distribution and configurable at a users local site. Software Support Tools ensures that software released by US-CMS is versioned and distributed using the CMS agreed upon software support tools: cvs and SCRAM.

4.3 Core Applications Software Budget and Personnel Requirements

Strategy for Estimating Resources

The resource profile for the CAS subproject is determined in a slightly different fashion from typical detector subprojects in which deliverables are agreed upon and then the costs are determined from a bottoms-up resource-loaded WBS. The method adopted for CAS acknowledges several differences in software projects compared to pure hardware projects.

- Firstly, for software there is no multi-year phase of mass-production of many similar components as there is for hardware. Each piece of the software is distinct from every other piece.
- Secondly, there is a continuous need for functional software systems from now to the turn-on of the experiment to support detector, trigger, and physics studies.
- Thirdly, the needs imposed on the software are continuously evolving, as are the software technologies which are exploited to meet those needs.

These factors together make it unrealistic to attempt to define highly detailed subtasks for the software subproject many years in advance.

This led us to adopt a "rolling planning" approach. Short to medium term tasks are defined in some detail with resource-loaded tasks and deliverables for US CMS in a similar fashion to a conventional hardware project. Tasks farther in the future, however, are defined in lesser detail and the resources are not explicitly estimated by rolling up low level tasks. Rather they are estimated on a level-of-effort basis by canonical scaling of the estimated International CMS needs, which are determined from an independent WBS for the complete CMS Software project. Then, as time progresses the future US CMS planning is refined such that the level of detail remains approximately constant as a function of time into the future.

Although the overall scope of the US CMS Applications Software subproject is bounded by this canonical scaling law, specific tasks are nonetheless assigned to US CMS individuals in a fully resource-loaded US-CMS WBS for detailed planning and tracking purposes. The resource estimates derived by rolling up the US-CMS WBS are then kept consistent with the canonical 25% scaling law for US contributions to CMS by adjusting the deliverables and schedule of US tasks. The level-of-detail for the US CMS CAS resource-loaded WBS is as follows:

- ***t* → *t* + 1 year** (i.e. **FY2001** for ***t*** = today)

Tasks are defined and resource-loaded to typically **level 5 / 6.**

- ***t* + 1 year → *t* + 2 years** (i.e. **FY2002** for ***t*** = today)

Tasks are defined and resource-loaded to typically **level 4 / 5**

- **beyond *t* + 2 years** (i.e. **FY2003, 04, 05,...**for ***t*** = today)

Tasks are defined to typically ~ **level 3,** often as an"ongoing resolution" philosophy.

This change process will be controlled by the project management, in consultation with US ASCB and CMS software and computing management. The process will be consistent with maintaining a fixed total cost and follow the change control procedures. The details of the change control procedure will need further discussion.

Overall constraints on the US CMS cost exposure need to be defined in agreement with all parties, including US CMS, CMS, the funding agencies and CERN. The nature of possible Memoranda of Understanding is under discussion in the Hoffmann Review of LHC Computing which will report to the CERN Resource Review Board in April 2001. Both the US-CMS and the CMS Software and Computing managements favor the more flexible procedure described above of defining level-of-effort contributions to CMS rather than inflexible software deliverables, since it makes optimal use of resources in a rapidly changing environment.

CAS Resource Requirements

The resource requirements of the US CMS CAS sub-project are the sum of the following components:

- **Professional Software Personnel for US CMS contributions to the CMS Core Software:** These resources are *constrained by definition to be 25% of the total CMS resource needs for all periods of the project.*
- **Professional Software Personnel resources for US-specific software support**: Each software engineer involved in the project is expected to spend approximately 1/4 of the time supporting US physicists (WBS item 2.4). This support load therefore *constitutes by definition 25% of the final US-CMS manpower total*, which is not used when calculating the 25% US-CMS contribution to International CMS.
- **Peripheral items to support software engineers**: A modest amount of support is required for each engineer to cover items such as desktop systems, travel, training, documentation, and software licenses (where not covered elsewhere by CMS or US CMS).
- **Management reserve:** There are a number of uncertainties associated to the CAS subproject. The resources are dominated by personnel with a wide variety of skills, more tracking experience is needed to understand how much effort various tasks really require, there is an intrinsic uncertainty in overall CMS estimate, and market forces influence salaries. While the adverse effects of the above may result in designing the software somewhat to cost, we consider it wise to give the project managers some ability to deal with unforeseen US CMS problems needing modest injections of manpower, perhaps expert consulting. This is set to 10% of the project cost, like for the UF subproject, and taken into account in the overall annual budget.

It should perhaps be emphasized that *no support* is included in the CAS subproject for any physics software work. The physicists working on software (roughly a factor of four more numerous than software professionals) are funded from off-project sources, typically the HEP base program.

The International CMS Software sub-project requires a total of 27 FTE's of software professionals in 2001 rising to a plateau of 33 FTE's in 2003. The incremental requirement is estimated to remain approximately flat until several years after 2005. This is to permit the shakedown and optimization of the operational system with a real and increasing data set and significantly increased use activities.

The corresponding FTE requirements for the US CMS Core Applications Software sub-project are shown in Table 2. The support personnel have been rounded to yield integer FTE's in the total requirement. The CAS subproject cost profile is obtained assuming the annual cost per engineer is on average \$158 k / FTE, with some variation for geographical location and skills. This includes overhead, fringe, desktop hardware, travel, training documentation, and software licenses.

The costs, uncorrected for inflation, as shown in Table 4, rise from \$1.58 M in fiscal year 2001 to \$2.05 M in 2005. The year 2006 is shown to indicate the required resources during each year of operation of the CMS experiment.

	2001	2002	2003	2004	2005	2006
CAS Project						
CAS FTE request	10	11	12	13	13	13
CAS Funding req in Million FY01$	1.58	1.74	1.90	2.05	2.05	2.05

Table 4: Estimated requirements for the US CMS Core Application Software subproject, including US-specific support and the cost for the sub-project, including personnel cost and associated support, such as travel, desktops, training and software licenses

5. Milestones and Budget for Overall Project

5.1 High Level Project Milestones

The high level milestones for the US CMS Software and Computing Project are chosen to be consistent with the overall high level milestones of the CMS Software Plan and the LHC/CMS schedule. These are shown in the sections for the User Facilities Subprojects and the Core Applications Software, respectively.

5.2 Project Office

This item provides personnel and operating funds to support the management of the project. It includes a staff of 1/2 FTE administrative assistant, 1/2 FTE Budget officer/planner, and 1 FTE assistant project manager/project engineer. The budget for travel for the L1PM and for support of reviewers, special consultants, etc is also under this item. The budget is currently absorbed in the User Facilities Subproject budget.

5.3 Cost and Labor Estimates

The requested resources for the CMS Software subproject and for the US CMS Core Applications Software subproject are estimated using the resource-loaded WBS. Costs are dominated by labor costs for software engineering manpower. Physicist manpower costs are not included.

The cost baseline will be established when it is approved by the Joint Oversight Group. The project cost baseline is equal to the sum of the budgeted costs for each element of the Work Breakdown Structure described above. Escalation is done assuming an annual escalation rate of 3% given as guidance from DOE. In addition an overall management reserve of 10% will be held each fiscal year and released during the year, as described above.

5.4 Budget Summary

Table 5 below lists the total budget for the US CMS Software and Computing Project. The table summarizes the resources needed for the UF and the CAS subprojects, plus the cost of the project office. The last row shows the total costs, escalated and including a management reserve of 10%. The User Facility budget contains the costs for Tier 2 hardware and personnel, which are shown separately.

FY	2001	2002	2003	2004	2005	2006	Total	2007
Project Office	**0.16**	**0.32**	**0.32**	**0.32**	**0.32**	**0.32**	**1.74**	**0.32**
CAS Total	**1.58**	**1.74**	**1.90**	**2.05**	**2.05**	**2.05**	**11.38**	**2.05**
UF Total	**2.41**	**3.48**	**4.17**	**8.89**	**10.74**	**11.43**	**41.13**	**10.64**
UF personnel	1.59	2.22	2.81	4.55	5.62	6.61	**23.39**	6.29
for Tier 1	*1.19*	*1.82*	*2.21*	*3.95*	*4.82*	*5.61*	***19.59***	*5.29*
for Tier 2	*0.40*	*0.40*	*0.60*	*0.60*	*0.80*	*1.00*	***3.80***	*1.00*
UF hardware	0.83	1.26	1.36	4.34	5.13	4.82	**17.74**	4.35
for Tier 1	*0.55*	*0.79*	*0.77*	*3.22*	*3.38*	*3.07*	***11.77***	*3.10*
for Tier 2	*0.28*	*0.47*	*0.59*	*1.12*	*1.75*	*1.75*	***5.96***	*1.25*
Total costs in Million FY01$	4.15	5.53	6.38	11.26	13.11	13.80	**54.24**	13.01
Total costs escalated, including managment res. [M$]	**4.57**	**6.27**	**7.44**	**13.54**	**16.24**	**17.60**	**65.65**	**17.09**

Table 5 Budget Summary of US CMS Software and Computing Project. Amounts are given in units of Million $ FY01, except for the last line. The costs shown for Tier2 centers in this table are for staff and hardware located at the Tier 2 centers.

6. Evolution of the US CMS Software and Computing Project

This plan applies to the ‘initial development and deployment phase’ for the software and facilities. Once that phase is complete, the software and the facilities go into the ‘operation, support, and further development phase’. This should occur a few years after CMS starts taking data. At that point, there should be a new ‘operations plan’ which would replace the project plan of the US CMS Software and Computing project. It is very important to recognize that software development and hardware evolution will continue throughout the life of CMS. The resources required for the ongoing operation of the facilities and evolution of the software and hardware are quite significant. The operation of the Regional Center is at least a 15-20 year commitment. For at least 2/3 of its lifetime, it will be in an ‘operations’ rather than ‘construction’ phase. Continual investment in software development, technology tracking, and R&D throughout this period will be essential if the facilities are to continue to serve the interests of US CMS as physics and computing technology move forward. Similarly, the scientific software will be in a state of continual development and evolution throughout the operations period. This will be driven both by changes in physics goals and analysis techniques and by changes in underlying software and hardware technologies.

REFERENCES

[1] US CMS User Facilities Subproject documentation, to be released
[2] US CMS Core Application Software Subproject documentation, to be released
[3] MONARC Project Execution Plan, LCB 98-003, September 1998; http://www.cern.ch/MONARC/

[4] US CMS Software and Computing Project Management Plan; http://home.fnal.gov/~bauerdic/uscms/scp/baserev_nov_2000/pmp.pdf
[5] CMS Computing Technical Proposal, December 1996, and reply of referees, April 1997. This schedule has recently been updated to reflect a more realistic LHC turn-on schedule and taking into account funding profiles.
[6] User Facility Sub-Project Hardware and Material Costs, to be released
[7] Hoffman Review Report, not yet public